RTFn: Enabling Cybersecurity Education through a Mobile Capture the Flag Client

Nicholas Capalbo, Theodore Reed, and Michael Arpaia
Computer Science Department, Stevens Institute of Technology, Hoboken, NJ, USA

Abstract— *Cybersecurity is one of the most highly researched and studied fields in computer science. It has made its way into numerous accredited universities as a full-fledged degree program. Students are constantly exposed to new technologies and methodologies through coursework. However there is a shortage of places to practice, in a controlled environment, the skills gained in the classroom. For that and numerous other reasons, the industry has seen a noticeable increase in capture-the-flag (CTF) style competitions. In the same vein, entering a CTF-like competition for the first time is a daunting task for any university. To aid in the organization of resources before, during, and after the competition we present the Rock The Flag network (RTFn). RTFn is a combination of hardware and software, which provides VPN capabilities as well as a central repository for tool tracking and real-time competition information. In this paper, we present an in depth discussion of this tool, its capabilities, and how it can aid in the organization required during CTF-style competitions.*

Keywords: Cybersecurity, Education, CTF, War-Games, Competition Logging, Collaboration

1. Introduction

Cybersecurity is an emerging concentration for undergraduate college students, and a developing concentration for masters candidates, doctoral, and post doctoral researchers. Academia has provided the field with a foundation for creativity, research, and innovation. As the field moves into the realm of undergraduate study, academia fails to adequately prepare students for the advanced technical work required at established security organizations [18]. Advanced degrees and programs build upon the experience of their candidates and students. Undergraduates often do not have such experience, thus it appears that academia is not an ideal venue for practical experience.

To compensate for a lack of experience, undergraduates must rely on programs which integrate that experience, such as labs or emulations. Alternatively, they can augment their undergraduate study with extra-curricular activities or internships. Many undergraduate programs recognize this requirement and integrate courses focusing on providing practical experience [21], [16], [3], [19]. These programs may invite experienced industry professionals [1], or allow students to use emulation and simulation sandboxes [12]. Bringing practical experience into the classroom is difficult.

It becomes even more challenging to provide red-team experience[1] to undergraduates. It is widely understood that this type of activity on commercially owned networks is against the law (as described in the Computer Fraud and Abuse Act of 1986). Although the goal of cybersecurity education and research is to create defensive strategies, playing the role of the attacker is often necessary. Defenses cannot be created effectively if attack methodologies are unknown. Creating attack taxonomies is the first step for assessing risks and developing defenses. This requirement elicits ethical considerations when providing students with the tools and processes for conducting attacks [1]. While these exercises demonstrate a great deal in a test bed environment, trust must be placed in students to not utilize these tools outside of a controlled environment.

Capture the Flag competitions are emerging as the solution to the issues created when introducing cybersecurity as a field for undergraduate study. These competitions [2], [6], [4], [15], [13] are designed to provide the needed cybersecurity experience that compensates current events, emerging trends, and course work. They provide attack and defense scenarios by facing students against difficult tasks, obscure procedures, and each other. Competitions provide feedback in the form of ranking and a detailed synopsis of the events. They introduce the fast-paced environment which surrounds the cybersecurity industry, and has the potential to teach information security related problem solving from experience [8].

In this paper we enhance this potential by providing institutions a quick solution to compete, perform well, collaborate amongst teams, identify weaknesses, and extract valuable experience during each event. We describe a hardware and software solution called the Rock the Flag network (RTFn). In section 2 we review work related to creating capture the flag compeitions, and their impact on collegiate study. Section 3 introduces our Rock the Flag network and provides an overview of its hardware and software suite. This section also outlines goals for capture the flag competition participation. In section 4 we examine the minimum set of hardware components for RTFn; section

[1]Experience that involves malicious behavior, such as penetration testing, or attack design.

5 examines the software components. In section 6 we outline our experiences as an institution introduced to capture the flag competitions and the steps we followed to organize a successful team of students. In sections 7 and 8 we provide our conclusions and future work.

2. Related Work

Capture the flag and cybersecurity competitions may utilize Virtual Private Networks (VPN) to enable competition play [10], [17], [2]. The VPN connects teams to either a defensive or offensive network where various oracles provide scoring mechanisms. In a defensive competition teams are often provided with a Virtual Machine (VM) which contains various flaws and security holes [18]. Teams may be scored by how well they can patch, secure, and defend their VM against a scoring system or other teams. If teams are required to defend their VM, as well as attack other team's VMs, the event is considered both and offensive and defensive competition. Competitions may also require only offensive challenges; these events typically involve a set of VMs maintained by the scoring system [14].

RTFn is a unique suite of hardware and software that enables collaboration. It resembles software-engineering and collaboration software, without focusing on development. RTFn does not suggest any methods to improve or change cybersecurity competitions. The tool is partly a response to documented "lessons learned" documents, published by various competition administrators, created to assist teams during competitions. RTFn may not be appropriate for all future competitions, and there are current competitions that will not actively utilize RTFn. We maintain that, in such competitions, RTFn will still enhance a team's performance during these competitions.

3. Approach

RTFn presents us with a significant amount of competition improvements and advancements. It is comprised of a combination of hardware and software that work with each other to maximize success in many of the areas of competition that often go overlooked. RTFn can be implemented as either a rack-mount solution or a mobile stand-alone system. Both implementations of RTFn have unique advantages and disadvantages. The rack-mount solution offers a static network address for persistent access by student competitors while the mobile stand-alone implementation offers added portability for off-site competitions.

RTFn was designed to solve the following existing problems related to participating in cybersecurty competitions:

- Universities may not have server space to host tools
- Teams may not have a dedicated meeting area to organize
- Teams experiences a lack of consistency and coherence between competition events
- It is difficult to extract learning items or recognize weaknesses during competition

RTFn was constructed using the following goals:

- Organize cybersecurity competition participation a-priori and post-priori
- Enable task-scheduling of competition challanges[2]
- Enable campus involvement, with minimum configuration and communication overhead
- Keep competition-related information secure
- Trend competition outcomes based on problem type, time, skills required, etc.

RTFn has very few requirements. There are a minimal amount of hardware requirements and the software used is adapted from open source solutions. Of the hardware requirements, it is necessary to equip RTFn with a large storage media. This is essential for the use of Virtual Machines, a repository of tools and the storage of data-mined information, challenges and reports. Also, it is imperative to outfit RTFn with a fast processor. This assists in many areas which include, but are not limited to, the running of virtual machines, the acceleration of key generation, running tools that support multi-threaded execution, and off-loading VPN requirements. It is important to keep the requirements to implement RTFn low to improve and promote university involvement in CTF-style competitions in an easy, fun fashion.

RTFn provides the following features: a challenge ownership portal, file uploading, and a real-time collaborative document editing. The challenge ownership portal assists in many areas such as work-load distribution, task completion, and coordination. Competitors can flag themselves as the owner of a specific challenge while they are working to minimize repetition. Marking ownership of a challenge is a highly effective way to accomplish multi-location participation. By marking their progress on a competition, team members will know to work on other challenges and propagate successful coordination and completion of challenges and tasks.

The RTFn also serves as a database of information security tools and scripts. Taking advantage of automated tools is an essential part of efficiently participating in CTF-stlye competitions. RTFn presents a structured way to categorize and cross-reference specific tools with specific types of challenges.

4. RTFn: Hardware Components

In this section we evaluate possible mobile hardware solutions, and provide a recommended configuration. We considered two deployment options for RTFn; one as a rack-mounted server in our university's information technology department's server room, and the second as a stand-alone

[2]We use the term competition challange throughout the paper. A challange may be a trivia question, deliverable, or acheivement. Typically these challanges are point scoring tasks.

mobile device. We chose to investigate the mobile device option for two reasons: rack-space may not be available to students at all universities, and a mobile option can include a mobile network which will quickly connect a physical lab- or group-environment. The mobile configuration we recommend also has the ability to be rack-mounted.

If students have access to rack-space then using a rack-mounted, dedicated metal, machine is the best option. Some of the software components described in the following section work best when they can be accessed before and after the competition. Whereas, a mobile device may change network addresses on campus and may not be accessible at all times, a stationary device will be able to be persistently reached at the same network address. A goal of RTFn is enabling university participation, thus rack-space is not a viable requirement. Note that RTFn should not be deployed as a virtual machine since it includes a hypervisor; it should be capable of running virtual machines for competitions. We strongly recommend hosting RTFn locally, on a university-owned network, since there is a possibility of logging offensive techniques used during competitions. Related offensive software should be stored for competition use only, labeled, and accessed securely.

RTFn has two hardware requirements: 1) a sufficiently large storage drive for storing competition-traffic captures, competition provided VMs, and associated collaboration data; 2) a CPU fast enough to run a virtual machine, maintain an OpenVPN connection, and generate keys to run a local OpenVPN. We also include optional features: a Wi-Fi B/G network interface, multiple Ethernet interfaces, and routing capabilities. The optional requirements enhance the mobile deployment option. We recommend using a Soekris computer to implement both the requirements and optional features. The Soekris net5501 [20] can be used as a stand-alone computer or racked with a special attachment.

5. RTFn: Software Components

In this section we describe the software components of the Rock The Flag network. The components can be divided into two groups: collaboration and reporting. We have identified that improvements to collaboration during cybersecurity competition will positively effect outcome. Based on related work, we also found that report generation, statistic tracking, and performance evaluation will also improve competition play. We describe each software component of RTFn as it relates to collaboration or reporting. Finally we conclude with a discussion of software security.

5.1 Collaboration

Robust project management software like Redmine [11], which combines wiki-style documentation and SVN support, present a particularly lucrative solution to information organization. During cybersecurity competition, however, time management is paramount. Wiki-style document editing, while used almost ubiquitously in information collaboration, is a time consuming process. Multiple participants may be simultaneously trying to update the same page, which may cause versioning conflicts. This type of information sharing also creates an overhead to those who are unfamiliar with the markup language syntax, which could cause time loss during the competition.

To solve this problem, RTFn implements a custom implementation of EtherPad [9], a web-based real-time collaborative document editor with chat support. Several additions are made to the EtherPad code base to include the following features, outlined in the following sections.

- Challenge ownership
- Related file uploading
- Meta-data labeling and challenge tagging

5.1.1 Challenge Ownership

To increase the overall efficiency of challenge completion during competitions, it is imperative for all participants to communicate and gain awareness of workload distribution. Without this awareness, competitors run the risk of duplicating already-completed work. Work distribution provides a sense of organization during the competition and allows competitors to focus their efforts intelligently. Competition challenges are often solved by multiple team players; unfortunately these challenges also suffer repeated work, which wastes a team's valuable time. Team members should not have to work through the same preliminary steps to solve a challenge. Furthermore, a second team member should be capable of picking up where another has finished by utilizing the collaborative document editor.

The challenge ownership feature is focused on improving team performance for competitions that last multiple days. Instead of requiring all students to be physically co-located, RTFn encourages distributed play. Coordination of tasks, and summaries of completed work, are provided by implementing ownership. Such that, if a student begins work on a challenge, they are assigned ownership; once they complete or exhaust their ability to continue the challenge, they can release ownership. This happens discretely, allowing students who play in different locations and at different times to keep their work and assignments synchronized.

RTFn's EtherPad implementation includes a dashboard of the challenges currently being attempted. Figure 1 shows a mock dashboard with 7 challanges; where (Ti) is the challenge title, (Tw) is the challenge type and (O) are the challenge owners. Each challenge shows a time counter, and allows an owner to mark the challenge as difficult or solved. When solved, the counter is paused. This dashboard identifies the "owners" of a particular challenge and allows other participants to quickly jump between questions. This mechanism, however, does not prevent participants from editing challenges being completed by other participants.

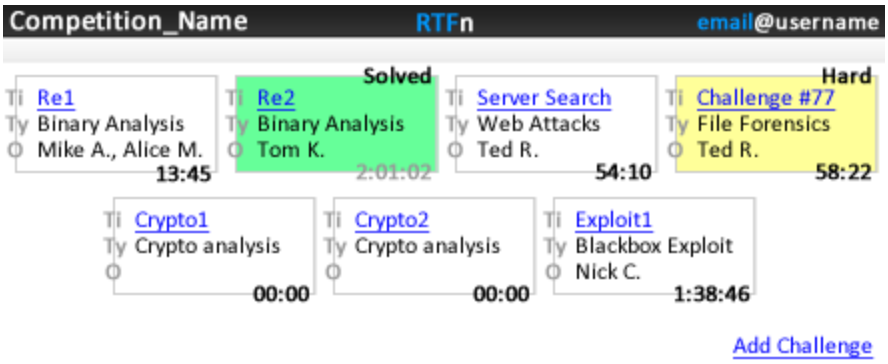

Fig. 1
A WEB VIEW OF THE CHALLANGE BOARD CREATED USING ETHERPAD.

In this scenario, competitors simply share ownership of the challenge, as shown in the first challenge.

5.1.2 Related file uploading

The original code base for EtherPad does not include functionality to support file uploading. We implement this feature to further aid in organization during the competition. It is very common to receive virtual machine images, PDFs, JPEGs, and other binary files during the competition for analysis and exploitation. To supplement the EtherPad code base, our file upload addition will be on a per-competition and per-challenge basis. That is, participants will have the ability to upload a file specifically related to whatever challenge they are attempting or for the more-general competition. When competitors switch between challenges, they will be able to view and download any files related to that particular challenge. When players view past competitions they'll be able to review general files, such as instructions, story-line documents, and team guides.

5.1.3 Meta-data labels and tags

Historical information retrieval provides insight into several facets of past competitions. Being able to search through past competitions, using data labeling, allows participants to observe past challenge strategies, tools, and general problem solving techniques. This serves as a great starting point for first-time competitors, reducing the lack of familiarity required to solve technically involved challenges. We describe these advantages in section 6. This form of data labeling is also used to trend focus areas of competition challenges. This methodology is discussed in the following section on report outline generation.

RTFn will supplement the EtherPad code base with support for meta-data labeling and tagging. After the completion of a competition, participants will have the opportunity to revisit challenges and tag them with labels. Labels are the higher level area of information security most closely associated with the problem. Some example labels could be web application, reverse engineering, code auditing, exploitation, trivia, password cracking, network forensics, cryptography, and social engineering. With appropriate labels in place, our implementation will support a further level of granularity using tagging. Tags will describe useful information regarding the methodology, tools, or related technologies required to complete or related to the challenge. Some example tags could be the related operating system, programming language, vulnerability language, associated CVE, as well as useful tools that were used to complete the challenge.

5.2 Reporting

Statistics gathering and report summary generation, performed after a competition, can be just as important and valuable as the preparation performed beforehand. Post-mortem report summary generation offers a holistic perspective into the competition flow [22]. Competition reports may highlight positive and negative actions and provide insight on future decision making. One of the problems RTFn attempts to solve is competition archival and inheritance. Many competitions are held annually, when competing multiple times RTFn can provide players with a refresher summary of past experiences. Report summaries may serve as an important tool for team and competition evaluation. Competitions may also require a report deliverable. RTFn can assist by providing an outline for writing. We implement two new features to EtherPad:

- Report outline generation
- Competition tracking

5.2.1 Report Outline Generation

Report outline generation provides the ability to trend competition challenges. By combining challenge outcomes with data labeling during collaboration, RTFn can provide insight on challenge types. A team can identify weaknesses by examining challenge types that are often not completed or

not attempted. A weakness might be a lack of specific subject understanding, experience, or time commitment. Identifying weaknesses is one of the primary goals of cybersecurity competition. This insight can help a team with recruitment, the structure of their practice, and perhaps the feedback provided to the university.

RTFn adds a reporting feature to the current EtherPad implementation, thus allowing competitors to analyze historical details of the competition. First, the reports will contain challenge-specific information such as the number and type of competition goals, with options for presenting higher and lower levels of granularity. After several competitions, teams will be able to trend goal patterns (e.g. offensive and defensive) to better predict and prepare for future iterations. Additionally, reporting will feature a timeline of events. Participants will be able to see a time breakdown of the competition, with a detailed view of how long each participant spent on each question and whether or not their effort led to a solution.

5.2.2 Competition Tracking

In order for a university's cybersecurity team to be successful, its participants need to constantly be aware of upcoming competitions. This awareness will allow for better planning and preparation during the weeks preceding the competition. With this knowledge, teams can accurately plan practices, scrimmages, and exercises to flesh out areas that may need more attention going into the competition.

RTFn recognizes the benefit of competition tracking; it implements a calendar add-on to EtherPad. With this functionality, teams will be able to mark important registration and competition dates, and plan their practices accordingly. The calendar framework will also send out email reminders to participants when an event nears. Competitors will also have the ability to export RTFn's calendar in iCalendar format to sync personal calendars.

The calendar feature's real power comes from a managed RSS feed of competitions. We maintain a list of security-related competitions and contests. Local RTFn installations can optionally poll this list and update their calendars. We imagine an interface for this public RSS feed that allows competition organizers to post their events. An example of the XML retreived from the managed RSS is shown in figure 2, this shows what information is contained for one competition. This should allow smaller competitions to gain popularity. However, the main goal of this managed RSS feed is to provide awareness to new competitors. This removes the overhead of discovering competitions, which aligns with RTFn's goal of removing organization overhead from competitions.

An example of the types of events that may appear in the feed may include:

- CSAW 2010 [15] - September 24-26th
- iCTF 2010 [2] - December 3rd

```
<title>East Coast Cyber-CTF</title>
<description>
  A security contest for high school
  and undergraduate students on the east coast.
  The first round contest will be held online at
  http://ecctf.example and the finals will be held
  in Washington D.C.</description>
<participation type=``remote'' method=``OpenVPN'' />
<time start=``May 4th, 2011 HH:MM:SS''
  end=``May 5th, 2011 HH:MM:SS'' />
<duration hours=24 />
<repeat annual=4 />
<pastinfo>
  <stats number_competitors=18 />
  <winner>Computer Security University</winner>
</pastinfo>
```

Fig. 2
EXAMPLE OF XML RETRIEVED FROM THE RSS FEED OF COMPETITIONS.

- Plaid CTF 2011 [4] - April 22-24th
- ISTS 2011 [16] - April 1-3rd
- ruCTF 2010 [10] - December 14th
- CODEGATE 2011 [5] - March 3-4th
- Defcon CTF [7] - Mid-Summer

5.3 System Security

Because of the offensive nature of CTF-style competitions and the capability for many of the stored tools to be used maliciously, it is imperative to securely store all aspects of the RTFn. RTFn's features work together during competition time to foster successful participation. The adapted EtherPad implementation, the repository of tools, the archive of old reports, the competition facing web servers, have a specific time that they are used in the realm of competition. The collaborative EtherPad software and competition facing web servers are used during competitions; the reports are used after the competition to analyze performance and before competitions to help prepare for future competitions. These features have a distinct purpose and it is important that they do not become leveraged for a counterproductive, malicious nature.

6. Campus Involvement

One of the most important features of RTFn is its ability to improve campus involvement and participation in CTF-style competitions. The many components of RTFn support the ease of getting involved in an extra-curricular cybersecurity organization and participating in competitions. The meta-data labeling and tagging system works to facilitate education and training. Also, students using RTFn for the first time are supplied with a wealth of organized data to browse and benefit from. The meta-data labeling and tagging system, as well as the file uploading system supplies students with a repository of challenges from previous competitions.

If a student wanted more ways to become involved in CTF-style competitions, they could browse the reports made by the collaborative document editing software. This will show new team members how specific problems were solved. This idea of community based self education removes the internal feeling of competition between team members by ridding the need of a team leader. It also promotes the idea of community based improvement and team building. The collaborative document editing software also facilitates task continuation and coordination with distant team members, making it easier for people to work together, regardless of where they might be located on campus.

Report outline generation and competition time lines offer students a detailed look into what competition is actually like. Students will be more comfortable with participating in a competition once they are more informed about the structure of it and will be able to create goals for themselves using the reports based on improving their weaknesses. The time line will assist in reducing the learning curve that often comes with CTF competitions by giving students a realistic expectation of time-based requirements.

In addition to making it easier for students to become involved in CTF-style competitions, RTFn also enables students to be more willing to become involved. Students feel more comfortable getting involved in a competition where there exists a plethora of information about what to expect. Often times, students are more comfortable participating from the comfort of their own dormitories, apartments, houses, etc., especially over the course of a several day competition. Given this, students will be inclined to take advantage of the collaborative document editing software features.

RTFn grants team leaders the ability to focus on the competition and assist other students instead of having to act as a systems administrator. The added organizational structure increases ease and fun, making students more inclined to participate. Since the software is open source and all materials are made available to everyone, fair team building is promoted. Also, since one of the biggest necessities of RTFn is collaboration, it is able to promote inclusion of newer team members and keep all students involved.

7. Conclusion

Early implementations of RTFn have been deployed on a home router and on a virtual machine running pfsense. From these deployments we created most of the hardware and software requirements outlined in this paper. Our team of students competed in CTF-related events for the first time and evaluated these deployments. Using RTFn, as described in this paper, enhances collaboration and productivity during CTF competitions. The hardware and software combination also enables a university to quickly gain a competitive advantage, record their performance, and evaluate their weaknesses. RTFn also provides valuable information to help improve undergraduate cybersecurity programs.

8. Future Work

To determine the effectiveness of our approach, we will rely on user experience data and performance monitoring. Observation during live competition is the most effective way to generate this data. This is accomplished by placing data monitors at different locations inside RTFn. We plan on offering RTFn to multiple universities participating in capture the flag competitions. To protect data privacy we will clearly explain what data will, and will not, be collected by RTFn. While the type of the data logged by RTFn is not sensitive in nature, participant awareness and disclosure is warranted.

We will monitor the performance of the various hardware components including the network interface(s), RAM, and CPU. For the network interface(s), we will monitor the total number of packets received versus dropped to help us gauge the amount of traffic routed during a competition. From this data, we can make improvements to RTFn to support a more reliable network interface card if required. In similar fashion, we can monitor CPU and RAM usage to better understand the processing load during peak competition involvement. To gauge the effectiveness of our customized EtherPad implementation, we will rely on user feedback; this interaction will help us better understand which are the most useful as well as possible component additions to be made.

We also plan to incorporate RTFn's modified EtherPad code base into a distributable disk image. Then, teams will not have to install the required software and troubleshoot any complications that may arise during the process. This will reduce the overhead for team organizers and ultimately foster more participation.

References

[1] AMAN, J. R., CONWAY, J. E., AND HARR, C. A capstone exercise for a cybersecurity course. *J. Comput. Small Coll. 25* (May 2010), 207–212.

[2] CHILDERS, N., BOE, B., CAVALLARO, L., CAVEDON, L., COVA, M., EGELE, M., AND VIGNA, G. Organizing large scale hacking competitions. In *Proceedings of the 7th international conference on Detection of intrusions and malware, and vulnerability assessment* (2010), DIMVA'10, pp. 132–152.

[3] CMU: CYLAB. Cylab: Confidence for a networked world. Website. `http://www.cylab.cmu.edu/`.

[4] CMU: PLAID PARLIAMENT OF PWNING. pCTF2011. Website. `http://www.plaidctf.com/`.

[5] CODEGATE. YUT Quals. Website. `http://yut.codegate.org/`.

[6] CONKLIN, A. Cyber defense competitions and information security education: An active learning solution for a capstone course. In *System Sciences, 2006. HICSS '06. Proceedings of the 39th Annual Hawaii International Conference on* (jan. 2006), vol. 9, p. 220b.

[7] DEF CON COMMUNICATIONS, INC. DEFCON, 2009. Website. `https://www.defcon.org/html/links/dc-ctf.html`.

[8] DODGE, R., HAY, B., AND NANCE, K. Standards-based cyber exercises. In *Availability, Reliability and Security, 2009. ARES '09. International Conference on* (March 2009), pp. 738–743.

[9] ETHERPAD. EtherPad Foundation, 2011. Website. http://etherpad.org.

[10] HACKERDOM. RuCTF, 2010. Website. http://www.ructf.org/.

[11] JEAN-PHILIPPE LANG. Redmine, 2011. Website. http://www.redmine.org/.

[12] LEE, C., ULUAGAC, A., FAIRBANKS, K., AND COPELAND, J. The design of netseclab: A small competition-based network security lab. *Education, IEEE Transactions on 54*, 1 (feb. 2011), 149 –155.

[13] LI, P., LI, C., AND MOHAMMED, T. Building a repository of network traffic captures for information assurance education. *J. Comput. Small Coll. 24* (January 2009), 99–105.

[14] MINK, M., AND FREILING, F. C. Is attack better than defense?: teaching information security the right way. In *Proceedings of the 3rd annual conference on Information security curriculum development* (2006), InfoSecCD '06, pp. 44–48.

[15] NYU POLY. Cyber Security Awareness Week. Website. http://www.poly.edu/csaw.

[16] NYU POLY: ISIS. The information systems and internet security (isis) laboratory. Website. http://isis.poly.edu.

[17] O'LEARY, M. A laboratory based capstone course in computer security for undergraduates. In *Proceedings of the 37th SIGCSE technical symposium on Computer science education* (2006), SIGCSE '06, pp. 2–6.

[18] POTHAMSETTY, V. Where security education is lacking. In *Proceedings of the 2nd annual conference on Information security curriculum development* (2005), InfoSecCD '05, pp. 54–58.

[19] RIT: SPARSA. Security practices and research student association. Website. http://www.sparsa.org/.

[20] SOEKRIS. net5501. Website. http://soekris.com/products/net5501.html.

[21] UCSB: SECLAB. The computer security group at ucsb. Website. http://www.cs.ucsb.edu/~seclab/.

[22] WAGNER, P. J., AND WUDI, J. M. Designing and implementing a cyberwar laboratory exercise for a computer security course. *SIGCSE Bull. 36* (March 2004), 402–406.